

**UTILITY
PATENT APPLICATION
TRANSMITTAL**
(Only for new nonprovisional applications under
37 C.F.R. 1.53(b))

Attorney Docket No. 1958.1031-000 (OID-1999-098-01)

First Named Inventor or Application Identifier Todd P. Guay

Express Mail Label No. EL 551548055US

jc922 U.S. Patent and
Trademark Office
11/01/00

Title of Invention

DATABASE INDEX VALIDATION MECHANISM

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. Fee Transmittal Form
(*Submit an original, and a duplicate for fee processing*)

2. Specification **[Total Pages 41]**
(*preferred arrangement set forth below*)

- Descriptive title of the invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to microfiche Appendix
- Background of the Invention
- Summary of the Invention
- Brief Description of the Drawings
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

3. Drawing(s) (35 U.S.C. 113) **[Total Sheets 1/24-24/24]**

- [] Formal Informal

4. Oath or Declaration/POA **[Total Pages 6]**

- a. Newly executed (original or copy)
- b. Copy from a prior application (37 C.F.R. 1.63(d))
(for continuation/divisional with Box 17 completed)
[NOTE Box 5 below]
 - i. DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).

5. Incorporation By Reference (*useable if Box 4b is checked*)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.

6. Microfiche Computer Program (*Appendix*)

7. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)

- a. Computer Readable Copy
- b. Paper Copy (identical to computer copy)
[] Pages
- c. Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. Assignment Papers (cover sheet & documents)

9. 37 C.F.R. 3.73(b) Statement Power of Attorney
(*when there is an assignee*)

10. English Translation Document (*if applicable*)

11. Information Disclosure Statement (IDS)/PTO-1449 Copies of IDS Citations

12. Preliminary Amendment

13. Return Receipt Postcard (MPEP 503)
(*Should be specifically itemized*)

14. Small Entity Statement(s) Statement filed in prior application, status still proper and desired

15. Certified Copy of Priority Document(s)
(*if foreign priority is claimed*)

16. Other: _____

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:[] Continuation Divisional Continuation-in-part (CIP) of prior application No.:

Prior application information: Examiner: Group Art Unit:

18. CORRESPONDENCE ADDRESS

NAME	Rodney D. Johnson, Esq. HAMILTON, BROOK, SMITH & REYNOLDS, P.C.			
ADDRESS	Two Militia Drive			
CITY	Lexington	STATE	MA	ZIP CODE
COUNTRY	USA	TELEPHONE	(781) 861-6240	FAX
	(781) 861-9540			

Signature	<i>Gerald M. Bluhm</i>	Date	11/11/00
Submitted by Typed or Printed Name	Gerald M. Bluhm	Reg. Number	44,035

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

FEE TRANSMITTAL FOR PATENT APPLICATIONS	Attorney Docket Number	1958.1031-000 (OID-1999-098-01)
	Application Number	
	First Named Inventor	Todd P. Guay

CLAIM CALCULATION (includes any preliminary amendment)

CLAIMS	(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) CALCULATIONS
	TOTAL CLAIMS (37 CFR 1.16(c) or (l))	54 - 20* =	34	x \$ 18 =	\$ 612
	INDEPENDENT CLAIMS (37 CFR 1.16(b) or (l))	4 - 3** =	1	x \$ 80 =	\$ 80
	MULTIPLE DEPENDENT CLAIMS (if applicable) (37 CFR 1.16(d))			+ \$ 270 =	\$
				BASIC FEE (37 CFR 1.16(a) or (h))	\$ 710
				Total of above Calculations =	\$ 1402
				Reduction by 50% for filing by small entity (37 CFR 1.9, 1.27, 1.28) =	\$
				TOTAL =	\$ 1402
				Surcharge - Late Filing of Declaration or Filing Fees (37 C.F.R. 1.16(e)) =	\$
				Petition for Extension of Time Fee (37 C.F.R. 1.17) =	\$
				Assignment Recordation Fee = (only when filed with application)	\$ 40
	* Reissue claims in excess of 20 and over original patent ** Reissue independent claims over original patent			TOTAL =	\$ 1442

1. Small entity status:

- a. A small entity statement is enclosed.
- b. A small entity statement was filed in the prior non-provisional application and such status is still proper and desired.
- c. Is no longer claimed.

2. A general authorization is hereby granted to charge deposit account number 08-0380 for any fees required under 37 CFR 1.16 and 1.17 in order to maintain pendency of this application. A copy of this authorization is enclosed for accounting purposes.

3. A check is enclosed for \$1,442. Please charge \$[] to Deposit Account No. 08-0380.

4. Other: _____

Signature	<i>Gerald M. Bluhm</i>	Date	11/11/00
Submitted by Typed or Printed Name	Gerald M. Bluhm, Esq.	Reg. Number	44,035

-1-

Date: 11-1-00

Express Mail Label No. EL 551548055US

Inventors: Todd P. Guay, Gregory S. Smith, Ari W. Mozes and
Gaylen D. Royal

Attorney's Docket No.: 1958.1031-000

DATABASE INDEX VALIDATION MECHANISM

BACKGROUND

Indexes are optional structures associated with database tables which exist primarily to enhance Structured Query Language (SQL) performance. Understanding and using indexes is important when optimizing SQL, because effective indexing can result in significant performance improvements in data retrieval.

The advantages of indexes do not come without a cost. Creating an index can be time consuming. In addition, indexes must be updated as the indexed table is modified by insertion or deletion of rows, or when an indexed table column is updated. As a result, indexes can degrade the performance of these operations. Furthermore, indexes consume additional disk space beyond that needed for the database.

Index tuning is the process of identifying an optimal set of indexes for each table in a database, given a set of SQL statements that are executed on the table, the properties of the table, and any restrictions imposed by the underlying database environment such as available disk space to create new indexes. Once a candidate set of indexes for the target table has been identified, it must be verified that the proposed solution is better than the current indexes created on the table.

This verification, however, involves changing the existing indexes to those of the proposed solution, and confirming that the optimizer uses the new indexes as expected and that no part of the workload is degraded in performance as a result of the index changes.

SEARCHED
INDEXED
SERIALIZED
FILED
10/24/00

15

20

Prior to the present invention, the verification process was typically performed by having an individual responsible for database administration, e.g., a database administrator (DBA), manually change or modify the indexes. A development database might be utilized during the verification process, to minimize any risks to the production database. A development database is a mirror copy of the production database and is used to test changes to the production database actual implementation. A primary advantage of using a development database is that risk is reduced, since the production database is not modified until the change has been confirmed on the development database. Major disadvantages of using a development database environment are that resources are required to maintain it, and that the DBA must ensure the development database accurately reflects the production database.

Verification typically involves running the set of SQL statements for the table, identifying any statements whose performance has degraded, performing a detailed evaluation to determine the cause of the degradation, and adjusting the index solution to resolve the degradation.

If a development database is not used, the verification is performed directly on the production database by making the recommended index changes and monitoring the production database for any problems that have been introduced.

SUMMARY

Indexing decisions are important and can be complex. Merely indexing everything is most likely not the best approach. Yet in order to achieve acceptable performance, indexes are generally used for databases. A solution is to index selectively, creating indexes only when the benefit outweighs the cost.

Given a workload, a current index set and a proposed index set, a Database Index Verification Mechanism can assist the user in determining which indexes should be created or modified to improve performance as well as identifying indexes that could be removed without degrading performance and without disabling an integrity constraint.

Therefore, an embodiment of the present invention can include a method and system that evaluates a plurality of candidate index sets for a workload of database statements in a database system by first generating baseline statistics for each statement in the workload. The workload can be reduced or collapsed into unique statements, and the statements can be SQL statements. An index superset can then be formed from a union of an existing or current index set and a proposed index set. The proposed index set may be provided, for example, by a user, or by an expert system. A candidate index set can be derived from the index superset, the candidate index being one of the plurality of candidate index sets. Selection of an index set may additionally be based on criteria which may include, but are not limited to, a cost value, a maximum number of allowed indexes and available storage space. Statistics can be generated based on the candidate index set and the baseline statistics, and presented along with statistics for the baseline and for the current index set, to a user or to an index tuning mechanism, for example.

The process of deriving a candidate index set and generating statistics based on the proposed index set can be executed repeatedly until at least one candidate index solution is found that adheres to user-imposed constraints and no further indexes can be removed from said candidate index solution without degrading performance of the workload. Furthermore, the current indexes can be disabled in order to establish the baseline. New candidate index solutions can be generated by eliminating at least one index within the candidate index solution that does not adhere to user-imposed constraints, which may be user-defined, such as a memory-usage constraint. In addition, new candidate index solutions can be generated by eliminating at least one index on a small table under evaluation if the index does not enforce an integrity constraint.

The statistics for a statement can be generated by first creating an execution plan that represents a series of steps to be taken by an optimizer in executing the statement. The statistics can include, but are not limited to, the number of executions of the statement, a user-defined weight of the statement, an index usage or a cost of the execution plan. The execution plan can then be evaluated, and statistics based on the

evaluation of the execution plan can be generated and recorded. The execution plan may be created without creating an index.

When determining the execution plan for a statement, the database's optimizer can examine available access paths, as well as statistics for the objects, such as tables or indexes, accessed by the statement. Evaluation of such an execution plan can include, for a table accessed by a statement under evaluation, identifying at least one index that would be used to retrieve data from the table upon an execution of the statement. Evaluation of such an execution plan can also include determining the cost of the execution plan. For example, the cost of the execution plan can be derived from the resource usage, such as CPU execution time or input/output access, required to execute the statement according to the execution plan.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of a Database Index Validation Mechanism, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a schematic illustration of a database management system embodying a database index validation mechanism for index tuning.

Fig. 2 is a schematic diagram illustrating the three stages of a particular database index validation mechanism.

Fig. 3A is a schematic diagram illustrating a high level overview of the baseline analysis stage of Fig. 2, in which a baseline cost is determined.

Fig. 3B is a block diagram illustrating a high level overview of the current index set analysis stage of Fig. 2.

Fig. 3C is a schematic diagram illustrating a high level overview of the proposed index solution analysis stage of Fig. 2.

Figs. 4A - 4C are a flowchart illustrating the steps performed by the workload evaluator of Figs. 3A-3C.

Fig. 5 is a schematic diagram illustrating a BaselineCost table as used by a particular Database Index Validation Mechanism.

Fig. 6 is a schematic diagram illustrating a IndexCost table as used by a particular Database Index Validation Mechanism.

Fig. 7 is a schematic diagram illustrating an IndexReference table as used by a particular Database Index Validation Mechanism.

Fig. 8 is a schematic diagram illustrating an IndexMaintenance table as used by a particular Database Index Validation Mechanism.

Fig. 9 is a flowchart illustrating the operation of the index evaluator of Figs. 3B and 3C.

Fig. 10 is a schematic diagram illustrating an IndexRefinement table as used by the index evaluator of Figs. 2 and 3 in a particular Database Index Validation Mechanism.

Fig. 11 is a flowchart illustrating an algorithm with which the cost reduction CR , as used in Equation 3, is calculated for a selected index.

Fig. 12 is a schematic diagram illustrating a SolutionRollup table as used by a particular Database Index Validation Mechanism.

Fig. 13 is a flowchart illustrating the steps executed by an embodiment of the index refiner of Fig. 3C.

Fig. 14 is a schematic diagram illustrating an exemplary set of four tables and a current index set comprising two indexes.

Fig. 15 is a diagram of a sample workload to be evaluated for the example of Fig. 14.

Fig. 16 is an illustration of an evaluation plan for the workload of Fig. 15 for a baseline solution as produced by a particular database index validation mechanism.

Fig. 17 is a schematic diagram illustrating how the BaselineCost table of Fig. 5 is populated for each workload statement of the workload of Fig. 15 and as a result of the evaluation of Fig. 16.

Fig. 18 is an illustration of an evaluation for the workload of Fig. 15 with current indexes enabled.

Fig. 19 is a schematic diagram illustrating how the IndexCost table of Fig. 6 is populated as a result of the evaluation of Fig. 18.

Fig. 20 is a schematic diagram illustrating the population of the IndexReference table of Fig. 7 as a result of the evaluation of Fig. 18.

Fig. 21 is a schematic diagram illustrating the population of the IndexMaintenance table of Fig. 8 as a result of the evaluation of Fig. 18.

Fig. 22 is a schematic diagram illustrating how the IndexRefinement table of Fig. 10 is populated by the index evaluator of Fig. 3C as a result of the population of the IndexCost, IndexReference and IndexMaintenance tables as shown in Figs. 19-21 respectively.

Fig. 23 is a schematic diagram illustrating how the SollutionRollup table of Fig. 12 is populated by the solution rollup evaluator of Fig. 3C as a result of the population of the IndexRefinement table as shown in Fig. 22.

Fig. 24 is an illustration of an evaluation for the workload of Fig. 15 for an exemplary proposed solution index set.

Figs. 25 - 27 are schematic diagrams illustrating how the IndexCost table, the IndexReference table, and the IndexMaintenance table are populated as a result of the new proposed index solution.

Fig. 28 is a schematic diagram illustrating the population of the IndexRefinement table for the example of Fig. 14.

Fig. 29 is a schematic diagram illustrating the population of the SolutionRollup table.

DETAILED DESCRIPTION

Fig. 1 is a schematic illustration of a database management system 210 embodying a database index validation mechanism for index tuning. The database management system 210 is shown to include a processor 212, an I/O device 214 and a system memory 216 connected by a system bus 218. Here the database is shown stored locally on storage disks 220A and 220B. It should be understood however, that the database can also be stored remotely on one or more devices (not shown).

System memory 216 is shown to include a database workload 222 and a reduced workload 4. The database workload 222 can reside in memory 216 as shown, within the database, or in any other secondary storage. Typically, a database workload 222 includes all database statements submitted to a database over a specified period of time. Several workload collection mechanisms can exist, three of which are discussed below.

A first mechanism filters application code, searching the source code for all database statements. A second mechanism introduces either a software or hardware collection process to the database environment between an application and the database. The process reads database statements as they are submitted to the database for processing and writes the statements to memory.

A third collection mechanism uses a database cache. Some database systems maintain an internal cache of frequently accessed database statements which have run against the database. This cache is typically used to share database statements between different users to save processing time. Since the shared database statements do not have to be reparsed, however, the cache may also be used to collect workload information. For ease of explanation and clarity of description, the remainder of the description assumes that the database workload is provided by a database cache capturing SQL statements.

System memory 216 is further shown to include an index validation system 224, which operates on a database workload 222 or a reduced workload 4. The index validation system 224 includes a workload evaluator 6, an index solution evaluator 9, and index solution rollup evaluator 12 and an index refiner 16.

In one embodiment, the workload evaluator 6 evaluates the reduced workload 4 for a first set of indexes for the workload. The index solution evaluator 9 determines a "value" for each index in the given index set. The solution rollup evaluator 12 rolls up the workload evaluation, providing a single cost metric for the first index solution.

5 Finally, the index refiner 16 creates a new candidate index solution by removing one or more indexes from the first solution. This new candidate index solution can be fed back through the index validation system.

10 The database workload 222 includes all SQL statements submitted to the database over a period of time. Typically, not all SQL statements included in the database workload are of interest to the index validation system 224. For example, many of the statements may define or modify the structure of the database or address the security for the database. These types of SQL statements generally do not address data retrieval efforts and are therefore not relevant to the index validation system 224. Therefore, providing the complete database workload 222 to the index validation system 224 may be inefficient and wasteful of processing resources.

15 Accordingly, a workload preprocessor (not shown), described in U.S. Serial No. 09/398,616, filed September 17, 1999, incorporated herein by reference, reduces a workload of a database system, such that the index validation system 224 can operate more efficiently. In addition, duplicate statements can be deleted, since evaluating the same statement several times would be wasteful. The result is a "reduced workload" 4.

20 The term "workload" is generally used below to refer to a reduced workload 4, although it should be recognized that a full workload 222 could also be employed.

Fig. 2 is a schematic diagram illustrating the three stages of a particular database index validation mechanism.

25 First, in the baseline stage 250, analysis is performed with all indexes off, except for those indexes that enforce constraints.

Next, a current index set analysis 252 is performed on the current index set 2, i.e., the set of indexes which have already been created for the database.

Finally, a proposed index set 3 is merged at 256 with the current index set 2 to form a “proposed index solution” 5, and a proposed index solution analysis 254 is performed, taking into account the results of the baseline analysis 250 and the current index set analysis 252. The result of the proposed index solution analysis 254 is a set of proposed refined index solutions 5A.

Fig. 3A is a schematic diagram illustrating a high level overview of the baseline analysis stage 250 of Fig. 2, in which a baseline cost is determined. The baseline analysis stage 250 records the performance of the workload 4 assuming that no indexes exist other than indexes that enforce integrity constraints. Thus, the current set of indexes on the tables under evaluation 7, excluding indexes which enforce integrity constraints, are disabled.

The current SQL workload 4 being executed against the tables 7 is fed into a workload evaluator 6. The workload evaluator 6 determines the cost of execution plans 10 for the individual statements of the workload 4 assuming that no indexes are available. This information is saved in a Baseline Cost table 80 for future comparisons in order to determine a proposed index’s “value” and thus serves as a baseline.

An execution plan is a description of the combination of steps to be performed by the database server to execute a SQL statement, such as SELECT, INSERT, UPDATE or DELETE. An execution plan includes an access method for each table to be accessed by the SQL statement, as well as an ordering of the tables, i.e., the join order. The “cost” of an execution plan is an estimated value which is proportional to the expected resource usage needed to execute the SQL statement with this execution plan. See, for example, “Oracle7 Server Concepts,” Chapter 13 (“The Optimizer”), Oracle Corporation, 1995, incorporated herein by reference.

Fig. 3B is a block diagram illustrating a high level overview of the current index set analysis stage 252 of Fig. 2.

Once the baseline costs have been determined, the current set of indexes 2 on the tables under evaluation 7, as well as the SQL workload 4 being executed against the tables 7, are fed into the workload evaluator 6. The workload evaluator 6 determines

the impact of the current index set 2 on the individual SQL workload statements 4. The cost of execution plans for the individual statements are recorded in the IndexCost table 90, while other relevant data are recorded in the IndexReference and IndexMaintenance tables, 100 and 110 respectively, as described in more detail below.

5 The index evaluator 9 determines a “value” for each individual index within the current index solution 2, and records the value in the IndexRefinement table 120.

These values are rolled up by the solution rollup evaluator 12 to determine the “efficiency” 14 of the current index solution 2, which is recorded in the SolutionRollup table 130.

10 Fig. 3C is a schematic diagram illustrating a high level overview of the proposed index solution analysis stage 254 of Fig. 2. Once the efficiency of the current index solution 2 has been determined and stored in the SolutionRollup table 130, a new set of indexes 3 is proposed by, for example, a user or a special tool that generates index recommendations. This proposed index set 3 is added to, or unioned with, the current index set 2 to form a proposed index superset 5. In one embodiment, the proposed indexes are evaluated by the optimizer as part of verification, but do not physically use space, nor are they visible to the optimizer outside of the verification process.

15 The workload evaluator 6 now determines the impact of the proposed index superset 5 by evaluating index usage for each SQL workload statement executing against the tables 7. A selector 11, which is part of the workload evaluator, selects a candidate index solution 5C from a proposed index solution set 5B.

20 As in the current index analysis stage of Fig. 3B, the index evaluator 9 determines the “value” each index provides, based on the data collected and stored in the BaselineCost table 80, the IndexCost table 90, the IndexReference table 100 and the IndexMaintenance table 110. This data is rolled up by the solution rollup evaluator 12 to determine the efficiency of the proposed index superset 5.

25 The solution refiner 16 refines the proposed index superset 5 by generating one or more refined index sets 5A that are subsets of the originally proposed index superset 5. Each of these refined index sets 5A is then fed back to the workload evaluator 6 and

the cycle continues until one of two conditions is true, as determined at 18. Together, the proposed index superset 5 and the refined index supersets 5A are considered “proposed index solutions” 5B.

One condition is that at least one index solution that adheres to user-imposed constraints exist and that no further indexes can be removed without degrading the performance of the SQL workload or disabling an integrity constraint.

The alternate condition that terminates the cycle is that all possible index solutions have been exhausted and none adhere to user-imposed constraints.

Once all viable index solutions have been analyzed, the index verification mechanism presents the individual index solutions at 20, ordered by their “efficiency,” to either the end user or an index tuning tool, for determining which index solution to implement. Further details are described below.

Workload Evaluator

Figs. 4A - 4C are a flowchart illustrating the steps performed by the workload evaluator 6 of Figs. 3A - 3C.

The primary goal of the workload evaluator 6 is to quantify the impact of a proposed solution 5, 5A, i.e., a set of indexes, on the individual SQL statements of the workload 4.

If the baseline is being generated, then first, in step 450, the workload 222 is collapsed, or reduced to unique SQL statements, i.e., a reduced workload 4, for example by deleting all duplicated statements and statements which do not manipulate data. As stated earlier, in one embodiment, this step is performed by a preprocessor.

Next, step 451 is executed, and any currently existing indexes 2 that do not enforce an integrity constraint are disabled. One method for accomplishing this is to modify the unique SQL statements associated with the tables under evaluation 7 by adding a hint that instructs the database server to ignore the specified index.

If, on the other hand, a current index set is being evaluated, the indexes are enabled at step 452.

Finally, if a proposed index solution 5B is being tested, then step 453 is executed, and a candidate index solution is selected from the proposed index solution 5B. The database server is then notified of the proposed indexes without actually creating them. One method for accomplishing this is to add a capability in the database server to allow a user to specify one or more index definitions of which the database server will be aware when generating an execution plan without requiring that the index actually exist.

Whether a current index set or a proposed index solution is being analyzed, tables 90, 100 and 110 are cleared, that is, all entries are removed (step 454).

For each unique SQL statement in the workload 4, the loop comprising steps 456-468 is executed. Step 456 determines whether there are any workload statements remaining to be analyzed.

If there are remaining statements, then in step 458, one of the remaining statements is selected. In step 460, an execution plan 10 is created based on available access paths and statistics for schema objects accessed by the selected SQL statement, i.e., tables and/or indexes. An optimizer (not shown) can generate the set of execution plans 3 for each workload statement based on available access paths and hints. This execution plan represents a series of steps used by the optimizer for executing the SQL statement.

In step 462, the execution plan 10 is evaluated. In the current index analysis stage 252, this includes identifying the index or indexes used to access rows of data from the database. The cost of the execution plan is then determined and stored in the BaselineCost table 80 for the baseline analysis stage 250 or the IndexCost table 90 for the current index or proposed index solution stages, 252, 254 respectively. The optimizer estimates the cost of each execution plan based on the data distribution and storage characteristics statistics for the tables and indexes, and any other objects in a data dictionary.

The "cost" is an estimated value proportional to the expected resource use needed to execute the statement using the execution plan. The optimizer calculates the

cost based on the estimated computer resources, including but not limited to I/O, CPU time and memory, that are required to execute the statement using the plan.

The workload evaluator 4 records the determined cost of the execution plan. If the baseline is being generated, the cost is recorded in the BaselineCost table 80 (step 464A). If a current index set or proposed index solution is being tested, the cost is recorded in the IndexCost table 90 (step 464B). These tables are described in more detail below.

In addition, if the baseline is being generated, the number of executions of each SQL statement and user-defined weight of the SQL statement are also recorded. In one embodiment, this information is stored in the BaselineCost table 80. Note that the number of executions and the user-defined weight for a given SQL statement remains constant. Thus, in one embodiment, this information is recorded just once during the baseline stage 250.

For baseline analysis, the loop is complete, and control returns to step 456. If there are no more SQL statements to analyze, the workload evaluator 6 is finished. Otherwise, at step 466, a determination is made as to whether the execution plan involves data access using any of the indexes being evaluated. If so, step 468 is executed, and the index or indexes used in the execution plan are recorded in the IndexReference table 100, along with the number of times each index was used in the execution plan.

In the current index set stage 252 or the proposed index solution stage 254, when there are no more SQL statements to process at step 456, steps 470 - 474 are executed for each index on the table or tables under evaluation 7. At step 472, an index is selected. At step 474, the estimated space required for index creation and the "volatility" of the index are determined and recorded in the IndexMaintenance table 110.

The various tables used in a particular embodiment of the present invention are now described. However, it would be understood by one skilled in the art that other table compositions or storage means are equally valid.

Fig. 5 is a schematic diagram illustrating the BaselineCost table 80 as used by a particular database index validating mechanism. The table 80 comprises five fields: StmtId 82, Cost 84, Executions 86, Weight 88 and UsageCost 89. For each row in the table 80, the StmtId field 82 holds an identifier which uniquely identifies a SQL statement. The Cost field 84 holds the cost of the execution plan for this SQL statement, as determined in step 462 and recorded in step 464A of Fig. 4.

The Executions field 86 holds the number of times the corresponding SQL statement is to be executed according to the workload 4.

The Weight field 88 holds a user-defined indicator of importance of this SQL statement relative to other statements in the workload 4. This allows a user to optimize the performance of his business's most important workload statements. In one embodiment, the weight defaults to 1, but can be changed to any value greater than 0.

The UsageCost field 89 holds the "usage cost" for the identified statement. This is calculated as

$$\text{UsageCost} = C_{\text{baseline}} * X * W \quad (\text{Eq. 1A})$$

where C_{baseline} is the execution plan cost (field 84), X is the number of executions (field 86) and W is the weight (field 88).

Fig. 6 is a schematic diagram illustrating the IndexCost table 90 as used by a particular Database Index Validation Mechanism. This table 90 is similar to the BaselineCost table 80 of Fig. 5, but is used to store non-baseline information, and comprises only three fields: StmtId 92, Cost 94 and UsageCost 96. As discussed previously, Executions and Weight fields are not required because this information is constant and is already stored in the BaselineCost table 80.

For each row in the IndexCost table 90, the StmtId field 92 holds an identifier which uniquely identifies a SQL statement. The Cost field 94 holds the cost of the execution plan for the identified SQL statement, determined in step 462 of Fig. 4.

The UsageCost field 96 holds the calculated cost for the identified statement. This is calculated as

$$\text{UsageCost} = C_{\text{index}} * X * W \quad (\text{Eq. 1B})$$

5

where C_{index} is the cost, from field 94, X is the number of executions, from field 86 of the BaselineCost table 80 and W is the weight, from field 88 of the BaselineCost table 80.

Referring again to Fig. 4, if the execution plan involves accessing data with one or more indexes, as determined in step 466, the index(es) used in the execution plan are recorded, in step 468, in the table IndexReference table 100. The number of times each index is referenced in the execution plan is also recorded.

Fig. 7 is a schematic diagram illustrating an IndexReference table 100 as used by a particular index validation mechanism, comprising three fields: StmtId 102, IndexName 104 and IndexRefCount 106. The StmtId field 102 holds an identifier which uniquely identifies a SQL statement. The IndexName field 104 holds the name of an index used as an access method in the execution plan for the identified SQL statement. The IndexRefCount 106 field holds the number of times this index is used in the execution plan for this SQL statement.

20

Fig. 8 is a schematic diagram illustrating an IndexMaintenance table 110 as used by a particular index validation mechanism. This table 110 comprises three fields: IndexName 112, RequiredSpace 114 and Volatility 116. The IndexName field 112 holds the name of an index. The RequiredSpace field 114 field holds the estimated number of kilobytes required to create named index. The Volatility field 116 holds the percentage of the workload that results in this index having to be updated. An index must be updated for each row inserted into the indexed table, each row deleted from the indexed table, and for each row that is updated in the indexed table where the indexed column is update. The volatility V is calculated as:

25

$$Volatility = 100 * \frac{P}{Q} \quad (\text{Eq. 2})$$

where P is the number of executions of the statement that result in an update to the index, and Q is the total number of executions of all statements.

5

Index Evaluator

The index evaluator 9 determines an index's "value" in the proposed index solution using the data stored by the workload evaluator 6 in the BaselineCost 80, IndexCost 90, IndexReference 100 and IndexMaintenance 110 tables. This value is used by the solution refiner 16 to identify indexes that can be removed without substantially degrading performance.

Fig. 9 is a flowchart illustrating the operation of the index evaluator 9 of Figs. 3B and 3C. In step 902, all existing entries in the IndexRefinement table 120, described below with respect to Fig. 10, are removed.

Step 904 determines whether there are any unprocessed indexes in the candidate index solution 5C. If not, the index evaluator is finished. Otherwise, at step 906, an index from the candidate index solution 5C is selected. In step 908, a "value" is determined for the selected index and stored in the IndexRefinement table 120. Step 908 is described in detail below with respect to Figs. 10 and 11.

Fig. 10 is a schematic diagram illustrating the IndexRefinement table 120 as used by a particular index validation mechanism. The IndexRefinement table 120 comprises two fields: IndexName 122 and UsageValue 124.

The IndexName field 122 holds the name of an index.

The UsageValue field 124 holds a "usage value" which is calculated for the named index. The value is 0 if the IndexRefCount field 106 of the IndexReference table 100 for this index is 0. Otherwise, the UsageValue is either 0 or "*usagevalue*," whichever is greater, where

10
15
20
25

$$usagevalue = K * \left(1 - \frac{CR}{TC} \right) * 100 - L * V \quad (\text{Eq. 3})$$

5 Here, K and L are user-defined constants which allow the user to indicate an importance factor for favoring data retrieval for frequently executed statements (high K) or minimal volatility (high L). In one embodiment, these constants, K and L, are defined once for the entire index solution refinement phase, and must be greater than or equal to 0.

10 CR , the cost reduction due to the index, represents the percentage of “data retrieval improvement” over the baseline that is a result of the index. Data retrieval improvement is indicated through a lower execution plan cost.

15 Fig. 11 is a flowchart illustrating an algorithm used to calculate the cost reduction CR due to a selected index.

20 First, at step 500, CR is initialized to 0.

25 The loop comprising steps 501 - 515 is repeated until there are no more SQL statements to be processed from the candidate index solution, as determined in step 501. In step 503, a SQL statement from the workload 4 is selected for processing.

30 Step 505 determines whether the selected SQL statement uses the index being evaluated. If the SQL statement does not use this index, then step 507 is executed, setting a temporary variable $cost$ to the value held in the UsageCost field 89 of the BaselineCost table 80. Execution then proceeds to step 515, discussed below.

35 If step 505 determines that the SQL statement does use this index, the flow goes instead to step 509, which determines the total number N of index uses in the execution plan for this SQL statement, i.e.,

$$N = \sum_{\text{SmtId}} \text{IndexReference.IndexRefCount} \quad (\text{Eq. 4})$$

Step 511 determines the number M of times the index being evaluated is used in the execution plan for the selected SQL statement. This is held in the IndexRefCount field 106 of the IndexReference table 100.

Step 513 determines a cost for the selected SQL statement as

$$5 \quad cost = \left(cost_{baseline} - M \frac{cost_{baseline} - cost_{index}}{N} \right) * X * W \quad (\text{Eq. 5})$$

where $cost_{baseline}$ is the cost (field 84) stored in the BaselineCost table 80, $cost_{index}$ is the cost (field 94) stored in the IndexCost table 90, and X and W are, respectively, the number of executions, from field 86 of the BaselineCost table 80 and the weight, from field 88 of the BaselineCost table 80, as defined previously, all for the corresponding SQL statement.

At step 515, CR is incremented by the amount stored in the temporary variable $cost$.

Referring back to Equation 3, TC represents the total cost for the baseline. This value is calculated as the sum of all UsageCost 89 values for all SQL statements.

15 V is the volatility, as calculated by Equation 2 and stored in the IndexMaintenance table 110.

Solution Rollup Evaluator

20 The solution rollup evaluator 12 determines the “efficiency” of a current or candidate index solution for the entire SQL workload 4. The information stored in the IndexCost 90, IndexMaintenance 110 and IndexRefinement 120 tables contains the raw data needed to populate a single entry in the SollutionRollupTable, described below, for a given index solution. In addition to the “efficiency,” various information including the cost of the index solution, the total space requirements for the indexes and the volatility of the index solution are recorded.

25 Index solutions adhering to user-imposed constraints are presented by the solution rollup evaluator 12, sorted by their “efficiency,” to the end user or the index

tuning tool. A user-imposed constraint may include conditions such as a maximum number of indexes allowed and/or memory or other storage usage limitations.

Fig. 12 is a schematic diagram illustrating a SolutionRollup table 130, as used by a particular index validation mechanism. This table 130 comprises five fields: 5 SolutionId 132, Cost 134, SpaceRequired 136, Weight 138 and UsageCost 139. The SolutionId field 102 holds an identifier which uniquely identifies the index solution. The Cost field 104 holds the average “weighted cost” for all SQL statements associated with the tables under evaluation 7. This can be calculated simply by summing the UsageCost fields 96 of the IndexCost table 90 for all SQL statements in the workload 4.

10 The SpaceRequired field 136 holds the number of kilobytes required to create the indexes used at least once. It can be calculated as the sum over the RequiredSpace fields 114 of the IndexMaintenance table 110.

The Volatility field 138 holds a percentage of references to the tables under evaluation 7 that result in any index having to be updated.

15 The Efficiency field 139 holds an indicator of the overall benefit of the proposed solution. It is calculated as

$$Efficiency = 100 * \left(1 - \frac{cost_{SolutionRollup}}{cost_{baselineTotal}} \right) - Volatility \quad (Eq. 6)$$

20 where $cost_{SolutionRollup}$ is from the Cost field 134 of the SolutionRollup table 130, and where $cost_{baselineTotal}$, which represents the total cost for the baseline, can be calculated as the sum of all UsageCost fields 89 from the BaselineCost table 80.

Solution Refiner

25 Fig. 13 is a flowchart illustrating the steps executed by an embodiment of the index solution refiner 16 of Fig. 3C.

The index solution refiner 16 refines the candidate index solution 5C by generating one or more new proposed index solutions 5A until at least one index

solution is found which adheres to user-imposed constraints and no further indexes can be removed without degrading the performance of the SQL workload 4 and without disabling an integrity constraint.

5 For example, in step 302, the solution refiner 16 eliminates any index or indexes, which do not enforce an integrity constraint, within the current proposed index solution that do not adhere to user-imposed constraints, such as a maximum number of columns allowed in an index.

DATA SHEET
10
15
20

In step 304, the solution refiner 16 eliminates any indexes on tables under evaluation 7 where the table is sufficiently small that an index would never be a desirable access method, provided that such an index does not enforce an integrity constraint. That is, in accessing data in small tables, a full table scan is always more efficient than using an index access method.

In step 306, the solution refiner 16 further eliminates any indexes within the proposed index solution that are never used, that do not enforce an integrity constraint and whose columns are not the prefix of another index which is being used. For example, if Idx1 is an index on columns Col1, Col2 and Col3 for some table, Idx2 is an index on column Col4, and Idx3 is an index on columns Col1 and Col2, and if Idx2 and Idx3 are not used, the solution refiner 16 can eliminate index Idx2. However, index Idx3 cannot be eliminated because its columns are the same as some of Idx1's columns. This is due to the fact that the Solution Refiner 16 may later eliminate Idx1, and as a result, Idx3 may then be used.

25 Step 308 determines whether or not the candidate index solution 5C adheres to user-imposed constraints. If it does not adhere to user-imposed constraints, then step 310 follows.

At step 310, one or more new proposed index solutions 5A are generated which represent a subset of the indexes in the current proposed index solution, taking into consideration which user-imposed constraint or constraints were violated.

For example, if the candidate index solution 5C violates a user-imposed constraint that defines the maximum number of indexes allowed, then the one or more

indexes that provide the least amount of value can be eliminated to construct a new refined index solution 5A that contains no more than the maximum number of indexes allowed.

As another example, if the candidate index solution violates a user-imposed constraint which defines space usage, then multiple proposed refined index solutions 5A can be constructed which include various combinations of indexes whose space usage requirements adhere to the user-imposed constraint. For each proposed refined index solution, the cycle of Fig. 3C is repeated, starting with the workload evaluator 6 and ending with the solution refiner 16.

As an example, assume a candidate index solution comprises indexes Idx1 - Idx4 with usage values and space requirements as shown in Table 1 below. Assume also that index solutions are restricted to a 32Kb space usage.

Index Name	Columns	Usage Value	Required Space
Idx1	Col1, Col2, Col3	100	32
Idx2	Col1, Col2	0	6
Idx3	Col1	150	16
Idx4	Col4	200	32

Table 1

In this example, the following proposed refined index solutions might be generated by the solution refiner 16:

- Proposed Refined Solution 1: {Idx4}
- Proposed Refined Solution 2: {Idx1}
- Proposed Refined Solution 3: {Idx2, Idx3}

If, at step 308, it is determined that the candidate index solution does adhere to user-imposed constraints, then steps 312 and 314 are performed instead of step 310.

Step 312 considers natural breaks or gaps in the usage values of each index in the candidate solution to create proposed solutions which include only those indexes providing the most benefit.

Step 314 constructs multiple index solutions that represent the different possible combinations of indexes excluding solutions that include an index that will not be used within that solution and whose columns are the prefix of another index which is being used, and excluding solutions that do not include at least one index for each "large" table under evaluation.

For example, assume that a candidate index solution comprises several indexes as shown in Table 2 below:

Index Name and Columns	Usage Value
I1 (column 1, column 2, column 3)	200
I2 (column 1, column 2)	0
I3 (column 4)	150

Table 2

In this case, the following proposed solutions might be generated:

- Proposed Solution 1: I1
- Proposed Solution 2: I2
- Proposed Solution 3: I3
- Proposed Solution 4: I1, I3
- Proposed Solution 5: I2, I3

A possible solution comprising (I1, I2) is excluded in step 314 because index I2, whose Usage Value is 0, will never be used as long as index I1 is part of the index solution.

EXAMPLE

An example of an embodiment of the present invention is now provided.

Fig. 14 is a schematic diagram illustrating an exemplary set of four tables and a current index set comprising two indexes. This example is used in the following

discussion. Table TABLE1 600 has five columns COL_1 601 through COL_5 605. Table TABLE2 610 has two columns COL_1 611 and COL_2 612. Table TABLE3 620 has three columns: COL_1 621, COL_2 622 and COL_3 623. Finally, table TABLE4 630 has two columns: COL_1 631 and COL_2 632.

5 The current index set for this example comprises two indexes: index IDX3_1 625 defined on column COL_1 621 of TABLE3 620 on; and index IDX3_2 626 defined on columns COL_1 621 and COL_3 623 of TABLE3 620.

10 Fig. 15 is a diagram of a sample workload 650 to be evaluated for the example of Fig. 14. Each of the various SQL statements 651-658, identified respectively as S1 - S8, accesses one or more columns from one or more of tables TABLE1 600 through TABLE4 630. For example, statement S1 651 returns the value stored in column COL_4 604 for each row of TABLE1 600 in which column COL_2 602 holds the value “1”.

15 For this example, assume that index solutions should weight equally index solutions that improve data access retrieval and index solutions that favor fewer updates to the index. In other words, the constant values defined above with respect to the index solution refiner 16 are $K = L = 1$.

BASELINE ANALYSIS STAGE

20 To create a baseline, the workload evaluator 6 first disables existing indexes that do not enforce an integrity constraint (step 452 of Fig. 4), i.e., indexes IDX3_1 and IDX3_2. The workload evaluator 6 then generates execution plans and determines and records execution plan costs of individual SQL workload statements for each statement (steps 456 - 464).

25 Fig. 16 is an illustration of an evaluation plan 660 for workload 650 by a particular database index validation mechanism for a baseline solution. For example, the evaluation 661 for SQL statement S1 651 indicates that a full scan of table TABLE1 will be performed upon execution of statement S1. That is, every row will be examined. Similarly, evaluations 662 - 668 are shown for statements S2 - S8 respectively.

Fig. 17 is a schematic diagram illustrating how the BaselineCost table 80 is populated, for each workload statement, with the execution plan cost 84, the number of executions 86 of the statement within the workload, the user-defined weight 88, and the usage cost 89 as calculated using Equation 1. The number of executions is determined by the workload evaluator 6 from statistics stored, for example, along with the actual SQL text for the workload statement. For this example, the numbers of Fig. 17 have been assumed. Similarly, it is assumed that the given weights have previously been established.

5

10

CURRENT INDEX SET ANALYSIS STAGE

Now that the baseline has been established, the indexes in the current index set 2, IDX_1 625 and IDX_2 626 in this example, are evaluated.

15

First, the workload evaluator 6 reenables the indexes. Execution plans are then generated for the workload statements and the execution plan costs of the individual SQL workload statements are determined.

20

Fig. 18 is an illustration of an evaluation 660A for workload 650 by a particular database index validation mechanism with the current index set of Fig. 14 enabled. Note that evaluations 661A - 666A, corresponding to workload statements S1 - S6 respectively, are identical to the baseline evaluations 661 - 666 of Fig. 16. However, evaluations 667A and 668A, corresponding to workload statements S7 and S8 respectively, now reflect use of the indexes. Note that the cost of each statement S7, S8 has been drastically reduced.

25

Fig. 19 is a schematic diagram illustrating how the IndexCost table 90 is now populated, for each workload statement, with the execution plan cost 94, and the usage cost 96 as calculated using Equation 2.

Fig. 20 is a schematic diagram illustrating the population of the IndexReference table 100. The IndexRefCount 106 field holds, for each statement using index IDX3_2, the number of times this index is used in the execution plan for that SQL statement. For each of statements S7 and S8, the IndexRefCount is 1.

NOTICE: This document contains neither recommendations nor conclusions of the Federal Energy Regulatory Commission. It is the product of an independent contractor.

Fig. 21 is a schematic diagram illustrating the population of the IndexMaintenance table 110 is populated. The space required for each index can be determined by various means which are beyond the scope of the present invention. The volatility for index IDX3_2 is calculated according to Eq. 2 as $(50 / 1260) * 100$ where there are 50 executions of statement S8, which causes an update to index IDX3_2, and where there are a total of 1260 executions for all statements S1-S8. Note that the volatility for index IDX3_1 is 0, because IDX3_1 is only on table TABLE3, column Col.1, which is not updated by statement S8.

Using the information stored in the BaselineCost 80, IndexCost 90, IndexReference 100 and IndexMaintenance 110 tables, the “value” of each individual index is now determined.

Fig. 22 is a schematic diagram illustrating how the IndexRefinement table 120 is populated. Using Equation 3, the usage value for index IDX3_2 is calculated as

$$\left(1 - \frac{CR}{TC}\right) * 100 - V = \left(1 - \frac{225,610}{304,810}\right) * 100 - 3.97 = 22.03$$

where CR , the cost reduction due to the index has been evaluated, in accordance with Fig. 11 and Equations 4 and 5, as

$$\begin{aligned} & 13,000 + 1,300 + 195,000 + 13,000 + 650 + 1,010 \\ & + \left(147 - \frac{147-3}{1}\right) * 500 * 1.0 + \left(147 - \frac{147-3}{1}\right) * 50 * 1.0 \\ & = 225,610 \end{aligned}$$

and where TC , the baseline total cost is $\sum(\text{BaselineCost} * \text{UsageCost}) = 304,810$.

Based on the information stored in the IndexCost 90, IndexReference 100, IndexMaintenance 110 and IndexRefinement 120 tables, the solution rollup evaluator 12 adds an entry comprising the “efficiency” of the candidate index solution to the

SolutionRollup table 130 after verifying this solution adheres to user-imposed constraints.

Fig. 23 is a schematic diagram illustrating how the SolutionRollup table 130 is populated.

5 The cost (field 134) is calculated as $\sum(\text{IndexCostUsage.Cost}) = 225,610$.

10 The required space (field 136) is calculated as $\sum(\text{IndexMaintenance.RequiredSpace}) = 14,080$.

15 The volatility (field 138) is calculated as $(50/1260) * 100 = 3.97$, where $\sum(\text{BaselineCost.Executions}) = 1260$ and a total of 50 executions of statement S8 result in at least one index being updated.

20 The efficiency (field 139) is calculated, according to Equation 6, as $(1 - (225610/304810)) * 100 = 22.03$.

PROPOSED INDEX SOLUTION ANALYSIS

Now, a new proposed index set is generated, either manually by a user, or by a special software tool, as mentioned previously. Assume for this example that the new proposed index set contains the following indexes:

- * IDX1_1 on TABLE1's column COL_2;
- * IDX1_2 on TABLE1's columns COL_2, COL_3 and COL_1;
- * IDX2_1 on TABLE2's column COL_1; and
- * IDX2_2 on TABLE2's column COL_2.

These proposed indexes, along with the indexes from the current index set, i.e., IDX3_1 and IDX3_2, form the proposed index superset 5 (Fig. 3C) to be considered as the first candidate index solution, Solution 1, by the database server.

25 Execution plans are then generated for this new candidate Solution 1, and execution plan costs of individual SQL workload statements are determined.

Fig. 24 is an illustration of an evaluation 660A for workload 650 by a particular database index validation mechanism for the proposed Solution 1 index set. As a result

of the new indexes, the costs of statements S1 - S4 and S6 have been greatly reduced, as indicated by the respective evaluations 661B - 664B and 666B.

Execution plan costs for each SQL statement, as well as index usage and maintenance information, are recorded in the IndexCost, IndexReference and IndexMaintenance tables, 80, 90 and 100 respectively.

Figs. 25 - 27 are schematic diagrams illustrating how the IndexCost table 90, the IndexReference table 100 and the IndexMaintenance table 110 are populated as a result of the new proposed index solution.

Volatility for each of indexes IDX1_1 and IDX1_2 is calculated as $(10/1260) * 100 = 0.79$, where there are 1,260 total statement executions and 10 statement executions for statement S5 result in index IDX1_1 and IDX1_2 being updated.

Volatility for index IDX2_1 is calculated as $(40/1260) * 100 = 3.17$ where there are 1,260 total statement executions, and 40 statement executions for statement S6 which result in index IDX2_1 being updated.

Volatility for IDX3_2 is calculated as $(50/1260) * 100 = 3.97$ where 50 statement executions for statement S8 result in index IDX3_1 being updated.

Using the information stored in IndexCost 90, IndexReference 100, IndexMaintenance 110 and BaselineCost 80 tables, the index evaluator 9 determines the "value" of the individual indexes.

Fig. 28 is a schematic diagram illustrating the population of the IndexRefinement table 120 with these new usage values.

The usage value for index IDX1_1 is calculated as follows. First, the cost reduction due to TABLE1 is

$$\begin{aligned}
 CR_{IDX1_1} &= \\
 &\left(260 - \left(\frac{260-4}{1} \right) * 1 * 50 * 1.0 \right) + 1,300 + 195,000 + 13,000 \\
 &+ 650 + 1,010 + 73,500 + 350 \\
 &= 292,010
 \end{aligned}$$

The total baseline cost is $TC = \sum(\text{BaselineCost}.\text{UsageCost}) = 304,810$. The usage value for $IDX1_1$ is thus

$$\left(1 - \frac{292,010}{304,810} \right) * 100 - 0.79 = 3.41$$

The usage value of index $IDX1_2$ is calculated similarly:

$$\begin{aligned}
 CR_{IDX1_2} &= \\
 &13,000 + \left(260 - \left(\frac{260-3}{1} \right) * 1 * 10 * 0.5 \right) + \left(780 - \left(\frac{780-9}{3} \right) * 3 * 500 * 0.5 \right) \\
 &+ \left(260 - \left(\frac{260-3}{2} \right) * 1 * 10 * 0.5 \right) + 650 + 1,010 + 35,000 + 7,350 \\
 &= 98,432
 \end{aligned}$$

The usage value for index $IDX1_2$ is therefore $\left(1 - \frac{98,432}{304,810} \right) * 100 - 0.79 = 66.91$

Similarly,

$$CR_{IDX2_1} =$$

$$13,000 + 1,300 + 195,000 + \left(260 - \left(\frac{160-3}{2} \right) * 1 * 100 * 0.5 \right) \\ + 650 + \left(101 - \left(\frac{101-3}{1} \right) * 1 * 40 * 0.25 \right) + 73,500 + 7,350 \\ = 291,487$$

The usage value of index IDX2_1 is $\left(1 - \frac{191,487}{304,810} \right) * 100 - 3.17 = 1.23$.

Finally,

$$CR_{IDX3_2} =$$

$$13,000 + 1,300 + 195,000 + 650 + 1,010 \\ + \left(147 - \left(\frac{147-3}{1} \right) * 1 * 500 * 1.0 \right) + \left(147 - \left(\frac{147-1}{1} \right) * 1 * 50 * 1.0 \right) \\ = 225,610$$

5

and the usage value of index IDX3_2 is $\left(1 - \frac{225,610}{304,810} \right) * 100 - 3.97 = 22.03$.

Using the information stored in the IndexCost 90, IndexReference 100 and IndexMaintenance 110 tables, the solution rollup evaluator 16 adds an entry in the SolutionRollup table for the “efficiency” of the candidate Solution 1 after verifying that the solution adheres to user-imposed constraints.

Fig. 29 is a schematic diagram illustrating the SolutionRollup table 120 as it is now populated.

The cost is calculated as $\sum(\text{IndexCost} \cdot \text{UsageCost})$.

The required space is $\sum(\text{IndexMaintenance} \cdot \text{RequiredSpace})$,

5 The volatility for Solution 1 is calculated as 1260 total statement executions of which 100 statement executions (S5, S6, S8) result in at least one index being updated: $(100/1260) * 100 = 7.9$.

The efficiency for Solution 1 is calculated, according to Equation 6 as
 $((1 - (4,945 / 304,810) * 100) - 7.9 = 90.5$

10 Solution Refinement Phase

Using the information stored in the IndexRefinement table 120, the solution refiner 16 refines the proposed index solution by generating one or more proposed index solutions.

15 First, any indexes within Solution 1 that do not adhere to user-imposed constraints; for example, a constraint that defines the maximum number of columns allowed in an index, are eliminated. For this example, assume all indexes adhere to user-imposed constraints.

20 Second, any index within the Index Solution that is never used, does not enforce an integrity constraint and whose columns are not the prefix of another index which is being used, is eliminated. In this example, IDX2_2 and IDX3_1 are never used. However, only IDX2_2 can be eliminated. IDX3_1 cannot be eliminated because it is a prefix of IDX3_2, which is used.

25 Table 3 below compares the UsageValues of the indexes. Natural breaks become apparent which define several groups of indexes.

Group	Index	UsageValue
1	IDX3_1	0
2	IDX1_1	3.14
	IDX2_1	1.23
3	IDX3_2	22.03
4	IDX1_2	66.9

Table 3

One possible proposal is to include only those indexes providing the most benefit; for example, a solution proposing IDX1_2 and IDX3_2.

Other possible proposals could comprise the different possible combinations of indexes excluding indexes those that are used but that encompass another index's columns. For example, indexes IDX1_2 and IDX3_3 have the highest usage values, yet because IDX1_1's indexed columns are also the prefix of the indexed columns for IDX1_2, it is likely that any index access satisfied by IDX1_2 will also be satisfied by IDX1_1.

Obviously, there is a chance that data retrieval performance may suffer by excluding the encompassing index, i.e., index IDX1_2. However, the performance degradation may be acceptable, especially considering space usage requirements will be less and the solution volatility may be smaller. The following refined solutions provide access to each of tables TABLE1, TABLE2 and TABLE3 through at least one index, but each refined solution exclude at least one index from the original proposed index solution 5:

Proposed Refined Solution: {IDX1_2, IDX2_1, IDX3_1}

Proposed Refined Solution: {IDX1_2, IDX2_1, IDX3_1, IDX3_2}

Proposed Refined Solution: {IDX1_1, IDX2_1, IDX3_2}

Proposed Refined Solution: {IDX1_1, IDX2_1, IDX3_1}

Proposed Refined Solution: {IDX1_1, IDX2_1, IDX3_1, IDX3_2}

Proposed Refined Solution: {IDX1_1, IDX1_2, IDX2_1, IDX3_2}

Proposed Refined Solution: {IDX1_1, IDX1_2, IDX2_1, IDX3_1}

5

These refined solutions 5A (Fig. 3C) are then passed back to the workload evaluator 6, one solution at a time, for a determination as to whether a newly proposed index solution may not in fact be the best one.

It will be apparent to those of ordinary skill in the art that methods involved in the present system for evaluating indexes may be embodied in a computer program product that includes a computer usable medium. For example, such a computer usable medium can include a readable memory device, such as a hard drive device, a CD-ROM, a DVD-ROM, or a computer diskette, having computer readable program code segments stored thereon. The computer readable medium can also include a communications or transmission medium, such as a bus or a communications link, either optical, wired, or wireless, having program code segments carried thereon as digital or analog data signals.

While this invention has been particularly shown and described with references to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the invention encompassed by the appended claims.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995
1000
1005
1010
1015
1020
1025
1030
1035
1040
1045
1050
1055
1060
1065
1070
1075
1080
1085
1090
1095
1100
1105
1110
1115
1120
1125
1130
1135
1140
1145
1150
1155
1160
1165
1170
1175
1180
1185
1190
1195
1200
1205
1210
1215
1220
1225
1230
1235
1240
1245
1250
1255
1260
1265
1270
1275
1280
1285
1290
1295
1300
1305
1310
1315
1320
1325
1330
1335
1340
1345
1350
1355
1360
1365
1370
1375
1380
1385
1390
1395
1400
1405
1410
1415
1420
1425
1430
1435
1440
1445
1450
1455
1460
1465
1470
1475
1480
1485
1490
1495
1500
1505
1510
1515
1520
1525
1530
1535
1540
1545
1550
1555
1560
1565
1570
1575
1580
1585
1590
1595
1600
1605
1610
1615
1620
1625
1630
1635
1640
1645
1650
1655
1660
1665
1670
1675
1680
1685
1690
1695
1700
1705
1710
1715
1720
1725
1730
1735
1740
1745
1750
1755
1760
1765
1770
1775
1780
1785
1790
1795
1800
1805
1810
1815
1820
1825
1830
1835
1840
1845
1850
1855
1860
1865
1870
1875
1880
1885
1890
1895
1900
1905
1910
1915
1920
1925
1930
1935
1940
1945
1950
1955
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
2025
2030
2035
2040
2045
2050
2055
2060
2065
2070
2075
2080
2085
2090
2095
2100
2105
2110
2115
2120
2125
2130
2135
2140
2145
2150
2155
2160
2165
2170
2175
2180
2185
2190
2195
2200
2205
2210
2215
2220
2225
2230
2235
2240
2245
2250
2255
2260
2265
2270
2275
2280
2285
2290
2295
2300
2305
2310
2315
2320
2325
2330
2335
2340
2345
2350
2355
2360
2365
2370
2375
2380
2385
2390
2395
2400
2405
2410
2415
2420
2425
2430
2435
2440
2445
2450
2455
2460
2465
2470
2475
2480
2485
2490
2495
2500
2505
2510
2515
2520
2525
2530
2535
2540
2545
2550
2555
2560
2565
2570
2575
2580
2585
2590
2595
2600
2605
2610
2615
2620
2625
2630
2635
2640
2645
2650
2655
2660
2665
2670
2675
2680
2685
2690
2695
2700
2705
2710
2715
2720
2725
2730
2735
2740
2745
2750
2755
2760
2765
2770
2775
2780
2785
2790
2795
2800
2805
2810
2815
2820
2825
2830
2835
2840
2845
2850
2855
2860
2865
2870
2875
2880
2885
2890
2895
2900
2905
2910
2915
2920
2925
2930
2935
2940
2945
2950
2955
2960
2965
2970
2975
2980
2985
2990
2995
3000
3005
3010
3015
3020
3025
3030
3035
3040
3045
3050
3055
3060
3065
3070
3075
3080
3085
3090
3095
3100
3105
3110
3115
3120
3125
3130
3135
3140
3145
3150
3155
3160
3165
3170
3175
3180
3185
3190
3195
3200
3205
3210
3215
3220
3225
3230
3235
3240
3245
3250
3255
3260
3265
3270
3275
3280
3285
3290
3295
3300
3305
3310
3315
3320
3325
3330
3335
3340
3345
3350
3355
3360
3365
3370
3375
3380
3385
3390
3395
3400
3405
3410
3415
3420
3425
3430
3435
3440
3445
3450
3455
3460
3465
3470
3475
3480
3485
3490
3495
3500
3505
3510
3515
3520
3525
3530
3535
3540
3545
3550
3555
3560
3565
3570
3575
3580
3585
3590
3595
3600
3605
3610
3615
3620
3625
3630
3635
3640
3645
3650
3655
3660
3665
3670
3675
3680
3685
3690
3695
3700
3705
3710
3715
3720
3725
3730
3735
3740
3745
3750
3755
3760
3765
3770
3775
3780
3785
3790
3795
3800
3805
3810
3815
3820
3825
3830
3835
3840
3845
3850
3855
3860
3865
3870
3875
3880
3885
3890
3895
3900
3905
3910
3915
3920
3925
3930
3935
3940
3945
3950
3955
3960
3965
3970
3975
3980
3985
3990
3995
4000
4005
4010
4015
4020
4025
4030
4035
4040
4045
4050
4055
4060
4065
4070
4075
4080
4085
4090
4095
4100
4105
4110
4115
4120
4125
4130
4135
4140
4145
4150
4155
4160
4165
4170
4175
4180
4185
4190
4195
4200
4205
4210
4215
4220
4225
4230
4235
4240
4245
4250
4255
4260
4265
4270
4275
4280
4285
4290
4295
4300
4305
4310
4315
4320
4325
4330
4335
4340
4345
4350
4355
4360
4365
4370
4375
4380
4385
4390
4395
4400
4405
4410
4415
4420
4425
4430
4435
4440
4445
4450
4455
4460
4465
4470
4475
4480
4485
4490
4495
4500
4505
4510
4515
4520
4525
4530
4535
4540
4545
4550
4555
4560
4565
4570
4575
4580
4585
4590
4595
4600
4605
4610
4615
4620
4625
4630
4635
4640
4645
4650
4655
4660
4665
4670
4675
4680
4685
4690
4695
4700
4705
4710
4715
4720
4725
4730
4735
4740
4745
4750
4755
4760
4765
4770
4775
4780
4785
4790
4795
4800
4805
4810
4815
4820
4825
4830
4835
4840
4845
4850
4855
4860
4865
4870
4875
4880
4885
4890
4895
4900
4905
4910
4915
4920
4925
4930
4935
4940
4945
4950
4955
4960
4965
4970
4975
4980
4985
4990
4995
5000
5005
5010
5015
5020
5025
5030
5035
5040
5045
5050
5055
5060
5065
5070
5075
5080
5085
5090
5095
5100
5105
5110
5115
5120
5125
5130
5135
5140
5145
5150
5155
5160
5165
5170
5175
5180
5185
5190
5195
5200
5205
5210
5215
5220
5225
5230
5235
5240
5245
5250
5255
5260
5265
5270
5275
5280
5285
5290
5295
5300
5305
5310
5315
5320
5325
5330
5335
5340
5345
5350
5355
5360
5365
5370
5375
5380
5385
5390
5395
5400
5405
5410
5415
5420
5425
5430
5435
5440
5445
5450
5455
5460
5465
5470
5475
5480
5485
5490
5495
5500
5505
5510
5515
5520
5525
5530
5535
5540
5545
5550
5555
5560
5565
5570
5575
5580
5585
5590
5595
5600
5605
5610
5615
5620
5625
5630
5635
5640
5645
5650
5655
5660
5665
5670
5675
5680
5685
5690
5695
5700
5705
5710
5715
5720
5725
5730
5735
5740
5745
5750
5755
5760
5765
5770
5775
5780
5785
5790
5795
5800
5805
5810
5815
5820
5825
5830
5835
5840
5845
5850
5855
5860
5865
5870
5875
5880
5885
5890
5895
5900
5905
5910
5915
5920
5925
5930
5935
5940
5945
5950
5955
5960
5965
5970
5975
5980
5985
5990
5995
6000
6005
6010
6015
6020
6025
6030
6035
6040
6045
6050
6055
6060
6065
6070
6075
6080
6085
6090
6095
6100
6105
6110
6115
6120
6125
6130
6135
6140
6145
6150
6155
6160
6165
6170
6175
6180
6185
6190
6195
6200
6205
6210
6215
6220
6225
6230
6235
6240
6245
6250
6255
6260
6265
6270
6275
6280
6285
6290
6295
6300
6305
6310
6315
6320
6325
6330
6335
6340
6345
6350
6355
6360
6365
6370
6375
6380
6385
6390
6395
6400
6405
6410
6415
6420
6425
6430
6435
6440
6445
6450
6455
6460
6465
6470
6475
6480
6485
6490
6495
6500
6505
6510
6515
6520
6525
6530
6535
6540
6545
6550
6555
6560
6565
6570
6575
6580
6585
6590
6595
6600
6605
6610
6615
6620
6625
6630
6635
6640
6645
6650
6655
6660
6665
6670
6675
6680
6685
6690
6695
6700
6705
6710
6715
6720
6725
6730
6735
6740
6745
6750
6755
6760
6765
6770
6775
6780
6785
6790
6795
6800
6805
6810
6815
6820
6825
6830
6835
6840
6845
6850
6855
6860
6865
6870
6875
6880
6885
6890
6895
6900
6905
6910
6915
6920
6925
6930
6935
6940
6945
6950
6955
6960
6965
6970
6975
6980
6985
6990
6995
7000
7005
7010
7015
7020
7025
7030
7035
7040
7045
7050
7055
7060
7065
7070
7075
7080
7085
7090
7095
7100
7105
7110
7115
7120
7125
7130
7135
7140
7145
7150
7155
7160
7165
7170
7175
7180
7185
7190
7195
7200
7205
7210
7215
7220
7225
7230
7235
7240
7245
7250
7255
7260
7265
7270
7275
7280
7285
7290
7295
7300
7305
7310
7315
7320
7325
7330
7335
7340
7345
7350
7355
7360
7365
7370
7375
7380
7385
7390
7395
7400
7405
7410
7415
7420
7425
7430
7435
7440
7445
7450
7455
7460
7465
7470
7475
7480
7485
7490
7495
7500
7505
7510
7515
7520
7525
7530
7535
7540
7545
7550
7555
7560
7565
7570
7575
7580
7585
7590
7595
7600
7605
7610
7615
7620
7625
7630
7635
7640
7645
7650
7655
7660
7665
7670
7675
7680
7685
7690
7695
7700
7705
7710
7715
7720
7725
7730
7735
7740
7745
7750
7755
7760
7765
7770
7775
7780
7785
7790
7795
7800
7805
7810
7815
7820
7825
7830
7835
7840
7845
7850
7855
7860
7865
7870
7875
7880
7885
7890
7895
7900
7905
7910
7915
7920
7925
7930
7935
7940
7945
7950
7955
7960
7965
7970
7975
7980
7985
7990
7995
8000
8005
8010
8015
8020
8025
8030
8035
8040
8045
8050
8055
8060
8065
8070
8075
8080
8085
8090
8095
8100
8105
8110
8115
8120
8125
8130
8135
8140
8145
8150
8155
8160
8165
8170
8175
8180
8185
8190
8195
8200
8205
8210
8215
8220
8225
8230
8235
8240
8245
8250
8255
8260
8265
8270
8275
8280
8285
8290
8295
8300
8305
8310
8315
8320
8325
8330
8335
8340
8345
8350
8355
8360
8365
8370
8375
8380
8385
8390
8395
8400
8405
8410
8415
8420
8425
8430
8435
8440
8445
8450
8455
8460
8465
8470
8475
8480
8485
8490
8495
8500
8505
8510
8515
8520
8525
8530
8535
8540
8545
8550
8555
8560
8565
8570
8575
8580
8585
8590
8595
8600
8605
8610
8615
8620
8625
8630
8635
8640
8645
8650
8655
8660
8665
8670
8675
8680
8685
8690
8695
8700
8705
8710
8715
8720
8725
8730
8735
8740
8745
8750
8755
8760
8765
8770
8775
8780
8785
8790
8795
8800
8805
8810
8815
8820
8825
8830
8835
8840
8845
8850
8855
8860
8865
8870
8875
8880
8885
8890
8895
8900
8905
8910
8915
8920
8925
8930
8935
8940
8945
8950
8955
8960
8965
8970
8975
8980
8985
8990
8995
9000
9005
9010
9015
9020
9025
9030
9035
9040
9045
9050
9055
9060
9065
9070
9075
9080
9085
9090
9095
9100
9105
9110

CLAIMS

What is claimed is:

5 1. A method for evaluating a plurality of candidate index sets for a workload of database statements in a database system, the method comprising:
 generating baseline statistics for each statement in the workload;
 forming an index superset from a union of a current index set and a proposed index set;
 deriving a candidate index set from the index superset, the candidate index being one of the plurality of candidate index sets;
 generating statistics based on the candidate index set and the baseline statistics; and
 presenting the generated statistics.

10 2. The method of Claim 1, further comprising:
 generating current index statistics for the workload responsive to the current index set, the presented generated statistics comprising the generated current index statistics.

15 20 3. The method of Claim 1, further comprising:
 repeatedly deriving a candidate index set and generating statistics based on the proposed index set.

25 4. The method of Claim 3, further comprising:
 terminating the repeated execution when at least one candidate index solution is found that adheres to user-imposed constraints and no further indexes can be removed from said candidate index solution without degrading performance of the workload and without disabling an integrity constraint.

5. The method of Claim 1, wherein deriving the baseline statistics comprises disabling current indexes.
6. The method of Claim 1, wherein generating statistics for a statement comprises:
 - creating an execution plan which represents a series of steps for executing the statement;
 - evaluating the execution plan;
 - generating and recording statistics based on the evaluation of the execution plan.
7. The method of Claim 6, wherein creating an execution plan is based on available access paths.
8. The method of Claim 6, wherein creating an execution plan is based on statistics for at least one schema object accessed by the statement.
9. The method of Claim 8 wherein the at least one schema object is a table.
10. The method of Claim 8 wherein the at least one schema object is an index.
11. The method of Claim 6, wherein evaluating the execution plan comprises:
 - for a table accessed by a statement under evaluation, identifying at least one index that would be used to retrieve data from the table upon an execution of the statement.
12. The method of Claim 6, wherein evaluating the execution plan comprises:
 - determining a cost of the execution plan.

13. The method of Claim 12, wherein the cost of the execution plan is derived from a resource use needed to execute the statement according to the execution plan.
14. The method of Claim 13, wherein the resource use includes CPU execution time.
15. The method of Claim 13, wherein the resource use includes input/output access.
16. The method of Claim 6, wherein the statistics include the number of executions of the statement.
17. The method of Claim 6, wherein the statistics include a user-defined importance of the statement.
18. The method of Claim 6, wherein the statistics include an index usage.
19. The method of Claim 6, wherein the statistics include a cost of the execution plan.
20. The method of Claim 1, wherein the statements are SQL statements.
21. The method of Claim 1, wherein the workload is reduced into unique statements.
22. The method of Claim 1, wherein deriving a candidate index set is responsive to a predetermined maximum number of allowed indexes.
23. The method of Claim 1, wherein deriving a candidate index set is responsive to available storage space.

5

24. The method of Claim 1, wherein the proposed index set is provided by a user.
25. The method of Claim 1, wherein the proposed index set is provided by an expert system.

10

26. The method of Claim 1, wherein an execution plan is created without creating indexes which are not in the current index set.

15

27. A system for evaluating a plurality of candidate index sets for a workload in a database system, the workload derived from a plurality of statements, the system comprising:

20

a workload evaluator which evaluates each statement within the workload;

an index solution evaluator which, responsive to the workload evaluator, evaluates each index in a candidate index set with respect to the workload, the candidate index solution being one of the plurality of candidate index sets;

a solution/rollup evaluator which, responsive to the index solution evaluator, evaluates the candidate index solution; and

a solution refiner which, responsive to the solution/rollup evaluator, generates at least one new candidate index solution.

25

28. The system of Claim 27, wherein the solution refiner generates at least one new candidate index solution by eliminating at least one index within the candidate index solution that does not adhere to user-imposed constraints.

29. The system of Claim 28, wherein the constraint is a user-defined constraint.

30. The system of Claim 28, wherein the constraint is a memory-usage constraint.

31. The system of Claim 27, wherein the solution refiner generates at least one new candidate index solution by eliminating at least one index on a small table under evaluation, and wherein the at least one index does not enforce an integrity constraint.

5

32. The system of Claim 27, wherein the evaluation created by the workload evaluator comprises an execution plan for each statement which represents a series of steps for executing the statement, wherein the workload evaluator evaluates the execution plan, and generates and records statistics based on the evaluation of the execution plan.

33. The system of Claim 32, wherein each execution plan is created based on available access paths.

34. The system of Claim 32, wherein each execution plan is created based on statistics for at least one schema object accessed by the statement.

35. The system of Claim 34 wherein the at least one schema object is a table.

20 36. The system of Claim 34 wherein the at least one schema object is an index.

37. The system of Claim 32, wherein the workload evaluator, for a table accessed by a statement under evaluation, identifies at least one index which would be used to retrieve data from the table upon an execution of the statement.

25 38. The system of Claim 32, wherein the workload evaluator determines a cost of the execution plan.

39. The system of Claim 38, wherein the cost of the execution plan is derived from a resource use needed to execute the statement according to the execution plan.

40. The system of Claim 39, wherein the resource use includes CPU execution time.

5

41. The system of Claim 39, wherein the resource use includes input/output access.

42. The system of Claim 32, wherein the statistics include the number of executions of the statement.

10

43. The system of Claim 32, wherein the statistics include a user-defined importance of the statement.

44. The system of Claim 32, wherein the statistics include an index usage.

15

45. The system of Claim 32, wherein the statistics include a cost of the execution plan.

46. The system of Claim 27, wherein the statements are SQL statements.

20

47. The system of Claim 27, wherein the workload is reduced into unique statements.

48. The system of Claim 27, wherein the solution refiner is responsive to a predetermined maximum number of allowed indexes.

25

49. The system of Claim 27, wherein the solution refiner is responsive to available storage space.

50. The system of Claim 27, wherein the proposed index set is provided by a user.

51. The system of Claim 27, wherein the proposed index set is provided by an expert system.

5
52. The system of Claim 27, wherein an execution plan is created without creating indexes which are not in the current index set.

10
53. A computer program product for evaluating a plurality of candidate index sets for a workload of database statements in a database system, the computer program product comprising a computer usable medium having computer readable code thereon, including program code which:

15
generates baseline statistics for each statement in the workload;
forms an index superset from a union of a current index set and a proposed index set;
repeatedly

20
derives a candidate index set from the index superset, the candidate index superset being one of the plurality of candidate index sets, and

25
generates statistics based on the candidate index set and the baseline statistics; and
presents the generated statistics.

54. A computer data signal embodied in a carrier wave for evaluating a plurality of candidate index sets for a workload of database statements in a database system, comprising:

program code for generating baseline statistics for each statement in the workload;

program code for forming an index superset from a union of a current index set and a proposed index set;

program code for repeatedly

deriving a candidate index set from the index superset,

the candidate index superset being one of the plurality of candidate index sets, and

generating statistics based on the candidate index set and seline statistics; and

program code for presenting the generated statistics.


DATABASE INDEX VALIDATION MECHANISM

ABSTRACT OF THE DISCLOSURE

5

A method evaluates a plurality of candidate index sets for a workload of database statements in a database system by first generating baseline statistics for each statement in the workload. An index superset is formed by combining an existing or current index set and a proposed index set. A candidate index set is derived from the index superset, the candidate index being one of the plurality of candidate index sets. Statistics for a statement are generated by first creating an execution plan which represents an efficient series of steps for executing the statement given the candidate index set. The execution plan is evaluated, and statistics based on the evaluation of the execution plan are generated and recorded. The cost of the execution plan is then determined and statistics are generated. Statistics for each candidate index set are rolled up and presented to a user or an index tuning mechanism.

10
9
8
7
6
5
15
4
3
2
1






Fig. 2




Fig. 3A




Fig. 3B




Fig. 3C




Fig. 4A




Fig. 48




Fig. 4C

BaselineCost Table

<u>StmtId</u>	<u>82</u>	<u>Cost</u>	<u>84</u>	<u>Executions</u>	<u>86</u>	<u>Weight</u>	<u>88</u>	<u>UsageCost</u>	<u>89</u>

Fig. 5

IndexCost Table

<u>StmtId</u>	<u>92</u>	<u>Cost</u>	<u>94</u>	<u>UsageCost</u>	<u>96</u>

Fig. 6

IndexReference Table

<u>StmtId</u>	<u>102</u>	<u>IndexName</u>	<u>104</u>	<u>IndexRefCount</u>	<u>106</u>

Fig. 7

IndexMaintenance Table

<u>IndexName</u> 112	<u>RequiredSpace</u> 114	<u>Volatility</u> 116

Fig. 8

IndexRefinement Table

<u>IndexName</u> 122	<u>UsageValue</u> 124

Fig. 10




Fig. 9




Fig. 11

130
/30
SolutionRollupTable

SolutionId 132	Cost 134	SpaceRequired 136	Volatility 138	Efficiency 139

Fig. 12




Fig. 13

TABLE1 600

COL 1	COL 2	COL 3	COL 4	COL 5
<u>601</u>	<u>602</u>	<u>603</u>	<u>604</u>	<u>605</u>

TABLE2 610

COL 1	COL 2
<u>611</u>	<u>612</u>

TABLE3 620

COL 1	COL 2	COL 3
<u>621</u>	<u>622</u>	<u>623</u>




TABLE4 630

COL 1	COL 2
<u>631</u>	<u>632</u>

Fig. 14

WORKLOAD 650

```
651- /*S1*/ select COL_4 from TABLE1 where COL_2 = 1;  
652- /*S2*/ select COL_2, COL_1 from TABLE 1 where COL_2 in (1,5,10);  
653 { /*S3*/ select COL_1 from TABLE1 where COL_2 = 1 and COL_3 = 2  
union all  
      select COL_1 from TABLE1 where COL_2 = 3 and COL_3 = 4  
union all  
      select COL_1 from TABLE1 where COL_2 = 5;  
654 { /*S4*/ select COL_1 from TABLE1 where COL_2 = (select COL_2 from TABLE2  
where COL_1 = 10);  
655- /*S5*/ insert into TABLE1 values (1, 2, 3, 4, 5);  
656- /*S6*/ update TABLE2 set COL_1 = 10 where COL_1 = 20;  
657- /*S7*/ select COL_3 from TABLE3 where COL_1 = 1;  
658- /*S8*/ update TABLE3 set COL_3 = "Oracle Corp" where COL_1<30;
```

Fig. 15

661	{	S1 (Cost 260) SELECT STATEMENT TABLE ACCESS (FULL) OF TABLE1
662	{	S2 (Cost 260) SELECT STATEMENT TABLE ACCESS (FULL) OF TABLE1
663	{	S3 (Cost 780) SELECT STATEMENT UNION-ALL TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE1
664	{	S4 (Cost 260) SELECT STATEMENT FILTER TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE2
665	{	S5 (Cost 260) INSERT STATEMENT INTO TABLE1
666	{	S6 (Cost 101) UPDATE STATEMENT UPDATE OF TABLE2 TABLE ACCESS (FULL) OF TABLE2
667	{	S7 (Cost 147) SELECT STATEMENT TABLE ACCESS (FULL) OF TABLE3
668	{	S8 (Cost 147) UPDATE STATEMENT UPDATE OF TABLE3 TABLE ACCESS (FULL) OF TABLE3

Fig. 16

BaselineCost Table 80

StmtId 82	Cost 84	Executions 86	Weight 88	UsageCost 89
S1	260	50	1.0	13000
S2	260	10	0.5	1300
S3	780	500	0.5	195000
S4	260	100	0.5	13000
S5	260	10	0.25	650
S6	101	40	0.25	1010
S7	147	500	1.0	73500
S8	147	50	1.0	7350

Fig. 17

WORKLOAD EVALUATION 660A

660A	{	S1 (Cost 26) SELECT STATEMENT TABLE ACCESS (FULL) OF TABLE1
662A	{	S2 (Cost 260) SELECT STATEMENT TABLE ACCESS (FULL) OF TABLE1
663A	{	S3 (Cost 780) SELECT STATEMENT UNION-ALL TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE1
664A	{	S4 (Cost 260) SELECT STATEMENT FILTER TABLE ACCESS (FULL) OF TABLE1 TABLE ACCESS (FULL) OF TABLE2
665A	{	S5 (Cost 260) INSERT STATEMENT INTO TABLE1
666A	{	S6 (Cost 101) UPDATE STATEMENT UPDATE OF TABLE2 TABLE ACCESS (FULL) OF TABLE2
667A	{	S7 (Cost 3) SELECT STATEMENT INDEX (RANGE SCAN) ACCESS (FULL) OF IDX3_2
668A	{	S8 (Cost 3) UPDATE STATEMENT UPDATE OF TABLE3 INDEX (RANGE SCAN) ACCESS (FULL) OF IDX3_2

Fig. 18

IndexCost Table 90		
StmtId 92	Cost 94	UsageCost 96
S1	260	13000
S2	260	1300
S3	780	195000
S4	260	13000
S5	260	650
S6	101	1010
S7	3	1500
S8	3	150

Fig. 19

IndexReference Table 100		
StmtId 102	IndexName 104	IndexRefCount 106
S7	IDX3_2	1
S8	IDX3_2	1

Fig. 20

IndexMaintenance Table 110		
Index Name 112	Required Space 114	Volatility 116
IDX3_1	5760	0
IDX3_2	8320	3.97

Fig. 21

IndexRefinement Table 120

IndexName 122	UsageValue 124
IDX3_1	0
IDX3_2	22.03

Fig. 22

SolutionRollup Table 130

SolutionId 132	Cost 134	RequiredSpace 136	Volatility 138	Efficiency 139
current	225610	14080	3.97	22.03

Fig. 23

WORKLOAD EVALUATION

660B

661B	S1 (Cost 4) SELECT STATEMENT TABLE ACCESS (BY INDEX ROWID) OF TABLE1 INDEX (RANGE SCAN) OF IDX1_1 (NON-UNIQUE)
662B	S2 (Cost 3) SELECT STATEMENT INLIST ITERATOR INDEX (RANGE SCAN) OF IDX1_2 (NON-UNIQUE)
663B	S3 (Cost 9) SELECT STATEMENT UNION-ALL INDEX (RANGE SCAN) OF IDX1_2 (NON-UNIQUE) INDEX (RANGE SCAN) OF IDX1_2 (NON-UNIQUE) INDEX (RANGE SCAN) OF IDX1_2 (NON-UNIQUE)
664B	S4 (Cost 3) SELECT STATEMENT INDEX (RANGE SCAN) OF IDX1_2 (NON-UNIQUE) TABLE ACCESS (BY INDEX ROWID) OF TABLE2 INDEX (RANGE SCAN) OF IDX2_1 (NON-UNIQUE)
665B	S5 (Cost 260) INSERT STATEMENT INTO TABLE1
666B	S6 (Cost 3) UPDATE STATEMENT UPDATE OF TABLE2 INDEX (RANGE SCAN) OF IDX2_1 (NON-UNIQUE)
667B	S7 (Cost 3) SELECT STATEMENT INDEX (RANGE SCAN) ACCESS (FULL) OF IDX3_2
668B	S8 (Cost 3) UPDATE STATEMENT UPDATE OF TABLE3 INDEX (RANGE SCAN) ACCESS (FULL) OF IDX3_2

Fig. 24

IndexCost Table 90

StmtId <u>92</u>	Cost <u>94</u>	UsageCost <u>96</u>
S1	4	200
S2	3	15
S3	9	2250
S4	3	150
S5	260	650
S6	3	30
S7	3	1500
S8	3	150

Fig. 25

IndexReference Table 100

StmtId <u>102</u>	IndexName <u>104</u>	IndexRefCount <u>106</u>
S1	IDX1_1	1
S2	IDX1_2	1
S3	IDX1_2	3
S4	IDX1_1	1
S4	IDX1_2	1
S6	IDX2_1	1
S7	IDX3_2	1
S8	IDX3_2	1

Fig. 26

IndexMaintenance Table 110

IndexName <u>112</u>	RequiredSpace <u>114</u>	Volatility <u>116</u>
IDX1_1	5760	0.79
IDX1_2	9344	0.79
IDX2_1	5760	3.17
IDX2_2	5760	0
IDX3_1	5760	0
IDX3_2	8320	3.97

Fig. 27

Index Refinement Table 120

IndexName <u>122</u>	UsageValue <u>124</u>
IDX1_1	3.14
IDX1_2	66.91
IDX2_1	1.23
IDX2_2	0
IDX3_1	0
IDX3_2	22.03

Fig. 28

SolutionRollup Table 130

SolutionId <u>132</u>	Cost <u>134</u>	RequiredSpace <u>136</u>	Volatility <u>138</u>	Efficiency <u>139</u>
current	225610	14080	3.97	22.03
Solution1	4945	40704	7.9	90.5

Fig. 29

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Declaration for Patent Application

As a named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated next to my name;

I believe I am the original, first and sole inventor (if only one name is listed) or an original, first and joint inventor (if plural names are listed in the signatory page(s) commencing at page 3 hereof) of the subject matter which is claimed and for which a patent is sought on the invention entitled

DATABASE INDEX VALIDATION MECHANISM

the specification of which (check one)

[X] is attached hereto.

[] was filed on _____ as United States Application

Number or PCT International Application No. _____

and was amended on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is known by me to be material to patentability as defined in 37 C.F.R. §1.56.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			Priority Not Claimed	Certified Copy Filed? YES	NO
(Number)	(Country)	(Day/Month/Year filed)	[]	[]	[]
(Number)	(Country)	(Day/Month/Year filed)	[]	[]	[]
(Number)	(Country)	(Day/Month/Year filed)	[]	[]	[]

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

(Application Number)	(Filing Date)
(Application Number)	(Filing Date)

I hereby claim the benefit under 35 U.S.C. 120 of any United States application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose information known by me to be material to patentability as defined in 37 C.F.R. 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)

As a named inventor, I hereby appoint the attorneys and/or agents associated with
Hamilton, Brook, Smith & Reynolds, P.C., Two Militia Drive, Lexington, Massachusetts 02421-4799
Customer No. 21005,

and _____
to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Please send correspondence to:

Customer No. **21005**
HAMILTON, BROOK, SMITH & REYNOLDS, P.C.
Two Militia Drive
Lexington, MA 02421-4799

or

Address as follows:

Direct telephone calls to: Rodney D. Johnson Telephone No.: 781-861-6240

Direct facsimiles to: Rodney D. Johnson Facsimile No.: 781-861-9540

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole

or first inventor Todd P. Guay

Inventor's Signature Todd P. Guay Date 10/30/00

Residence 10 North London Drive

Nashua, New Hampshire 03062

Citizenship USA

Post Office Address (same as above)

Full name of second joint

inventor, if any Gregory S. Smith

Inventor's Signature Henry J. Smith Date 10/25/2000

Residence 16 Abbey Road

Merrimack, New Hampshire 03054

Citizenship **USA** _____

Post Office Address (same as above) _____

Digitized by srujanika@gmail.com

Full name of third joint

inventor, if any Ari W. Mozes

Inventor's Signature _____ Date _____

Residence 757 Elm Street #12

San Carlos, California 94070

Citizenship USA

Post Office Address _____

Full name of fourth joint

inventor, if any Gaylen D. Royal

Inventor's Signature Gaylen D. Royal Date 10/23/2000

Residence 47 Pembroke Way

Bedford, New Hampshire 03110

Citizenship USA

Post Office Address (same as above)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Declaration for Patent Application

As a named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated next to my name;

I believe I am the original, first and sole inventor (if only one name is listed) or an original, first and joint inventor (if plural names are listed in the signatory page(s) commencing at page 3 hereof) of the subject matter which is claimed and for which a patent is sought on the invention entitled

DATABASE INDEX VALIDATION MECHANISM

the specification of which (check one)

[X] is attached hereto.

[] was filed on _____ as United States Application

Number or PCT International Application No. _____

and was amended on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is known by me to be material to patentability as defined in 37 C.F.R. §1.56.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			<u>Priority Not Claimed</u>	<u>Certified Copy Filed?</u>
<u>YES</u>	<u>NO</u>			
(Number)	(Country)	(Day/Month/Year filed)	[]	[] []
(Number)	(Country)	(Day/Month/Year filed)	[]	[] []
(Number)	(Country)	(Day/Month/Year filed)	[]	[] []

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

<u>(Application Number)</u>	<u>(Filing Date)</u>
<u>(Application Number)</u>	<u>(Filing Date)</u>

I hereby claim the benefit under 35 U.S.C. 120 of any United States application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose information known by me to be material to patentability as defined in 37 C.F.R. 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)
(Application Serial No.)	(Filing date)	(Status: patented, pending, abandoned)

As a named inventor, I hereby appoint the attorneys and/or agents associated with
Hamilton, Brook, Smith & Reynolds, P.C., Two Militia Drive, Lexington, Massachusetts 02421-4799
Customer No. 21005,

and _____

to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Please send correspondence to:

Customer No. **21005**
HAMILTON, BROOK, SMITH & REYNOLDS, P.C.
Two Militia Drive
Lexington, MA 02421-4799

or

Address as follows:

Direct telephone calls to: Rodney D. Johnson Telephone No.: 781-861-6240

Direct facsimiles to: Rodney D. Johnson Facsimile No.: 781-861-9540

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole
or first inventor Todd P. Guay

Inventor's Signature _____ Date _____

Residence 10 North London Drive
Nashua, New Hampshire 03062

Citizenship USA

Post Office Address (same as above)

Full name of second joint

inventor, if any Gregory S. Smith

Inventor's Signature _____ Date _____

Residence 16 Abbey Road

Merrimack, New Hampshire 03054

Citizenship USA

Post Office Address (same as above)

Full name of third joint

inventor, if any Ari W. Mozes

Inventor's Signature Ari W. Mozes Date 10/27/00

Residence 757 Elm Street #12

San Carlos, California 94070

Citizenship USA

Post Office Address (same as above)

Full name of fourth joint

inventor, if any Gaylen D. Royal

Inventor's Signature _____ Date _____

Residence 47 Pembroke Way

Bedford, New Hampshire 03110

Citizenship USA

Post Office Address (same as above)